On Climate Agreements with Asymmetric Countries: Theory and Experimental Results

Charles F. Mason†
H.A. True Chair in Petroleum and Natural Gas Economics
Department of Economics
University of Wyoming,
1000 E. University Ave.
Laramie, WY 82071

June 21, 2019

Abstract

I model International climate agreements among asymmetric countries, each of whom must select a profile of CO_2 emissions over time. Predictions from this model imply larger reductions by "large" countries, but larger proportional reductions by "small" countries. I then analyze experimental data that sheds light on this issue. In contrast to the theoretical predictions, I find that smaller countries do not reduce emissions proportionately to their Nash level, and so the burden falls mostly on larger countries. Moreover, combined emissions are indistinguishable from the one-shot Nash emissions. This pessimistic outcome extends the commonly-found result in the literature that negotiations in similar repeated games (but with symmetric players) generally do not offer much hope for meaningful agreements, unless the effects are modest.

Keywords: Climate Negotiations, Repeated Game, Experiments

JEL Areas: D8, L15

†phone: +1-307-766-2178; fax: +1-307-766-5090; e-mail: bambuzlr@uwyo.edu

1 Introduction

Climate change is a difficult problem to address because it is global, requiring concerted action by multiple sovereign countries (Barrett, 2008). National sovereignty implies that international agreements require actions be in each signatory country's national self-interest. This feature has lead a number of scholars to conclude that such agreements are either implausible or ineffective (Barrett, 1994; Carraro and Siniscalco, 1993; Finus, 2001).

These problems are evident in efforts to negotiate an international agreement on climate change. Two significant agreements have emerged from international negotiations: the Kyoto Protocol in 1997, and the Paris Agreement in 2015. Among the features distinguishing these agreements, one stands out: the Kyoto Protocol entailed emissions reductions from developed countries, while the Paris Agreement on Climate Change involved all countries. Accordingly, participants in the Kyoto agreement were relatively similar in size and their level of economic development, while participants in the Paris Agreement were diverse.

The pessimistic message from the literature described above involves symmetric players; allowing for asymmetric countries seems likely to further undercut the efficacy of forming an international agreement. I explore this aspect in this paper. To this end, I analyze data from an experimental analysis based on a relatively simple linear-quadratic payoff structure, where players are of two types: "large countries" and "small countries" (for expositional convenience, I often refer to these types as "large players" and "small players" in the pursuant discussion). Large countries have payoff functions that yield larger rewards; this can be interpreted as resulting from lower abatement costs or lower marginal damages from emissions. In this way, large countries are emblematic of developed countries, while small countries represent developing countries. I find that subject choices in the presence of such heterogeneous payoffs are significantly closer to the (non-cooperative) Nash equilibrium than are choices in a symmetric structure. Moreover, the observed market shares for small countries exceed the shares such countries obtain in the Nash equilibrium.

This experimental outcome cannot be explained by standard theoretical treatments, as those generally allocate a larger share of combined emissions to the larger player.[1] But it is at least broadly consistent with features of the two climate treaties discussed above: In the Kyoto Protocol, large countries bore the load of curtailing emissions, suggesting that small countries – whom one might anticipate would best-respond to large countries' behavior – would likely increase their emissions. Under the Paris Agreement, countries are free to suggest emission reductions (through so-called "Intended Nationally Determined Contributions"). One concern expressed by the current administration of the United States (US) is that smaller countries would "take advantage of the US," presumably by raising their emissions in response to US reductions. The finding is also broadly consistent with casual empirical evidence taken from the numerous attempts – via the annual "Conference of the Parties" that have taken place since 1995 – to negotiate an international agreement limiting carbon emissions. For a long time, smaller (underdeveloped) countries insisted that larger (developed) countries bear most of the brunt in reducing emissions.

The paper is organized as follows. Section 2 presents a theoretical framework for cooperative arrangements in the climate negotiation game. I base this discussion on the "grim strategy," under which a cooperative emissions profile is undertaken so long as all players have honored the

[1] See, for example, Schmalensee (1987), who studies asymmetric industrial structures. In that setting, his "low cost firm" is analogous to the large player in my setting, and the "high cost firm" is analogous to the small player. His Table 1 shows that the increased payoff earned by the larger cost firm is smaller than would obtain under "proportional reduction" – in which case the firms' shares would correspond to the Nash levels. The result is consistent with the Equity, Reciprocity and Competition (ERC) model (Bolton and Ockenfels, 2000, p. 181).

agreement in the past, but where defection triggers a regime shift to the Nash equilibrium forever after. The ability of this strategy to deliver improvements, via reduced emissions, requires that two incentive constraints be satisfied – one for each player. I argue that the incentive constraint for the larger player is the more likely to bind, thereby requiring the arrangement to be disproportionately more favorable to the large player. In particular, such a regime would afford the large player a larger share of global emissions than would obtain under a pro-rata sharing rule; this fact makes agreement less probable. I then investigate the empirical plausibility of these predictions, using data from a battery of experiments. Section 3 offers a discussion of that experimental structure, while section 4 offers empirical results based on econometric analysis of the data. Section 5 offers discussion.

2 Modeling Cooperation Among Asymmetric Countries

To flesh out the theoretical backdrop to the experimental design discussed below, I consider a simple interaction involving two countries $k = 1,2$ each of whom must select emissions e_k; this design abstracts from carbon stock effects so as to sharpen the focus, and highlight the difficulties that arise from asymmetries. Net benefits from a combination of emissions depend on indirect benefits from emissions (for example, via increased economic activity) and direct damages, which depend on combined emissions $E = e_1 + e_2$. I assume the marginal damages from combined emissions are dE – and so are identical – for the two countries; damages are $D_k = dEe_k$. Asymmetries arise from potentially different net benefits. This difference could reflect a more advanced economy for one of the countries, as when one is developed and the other underdeveloped or developing; alternatively, it might reflect some sort of abatement cost advantage for the larger country, perhaps because of better institutions or superior technological capabilities. To fix ideas, I let country 1 be a large country. In a symmetric structure, country 2 is also large, while in an asymmetric structure country 2 is small. The marginal benefit of emissions for country i is b_i, where my notational convention implies $b_1 \geq b_2$, with strict inequality in the asymmetric game. Payoffs for country i are then

$$\pi_i = b_i e_i - dEe_i \equiv (b_i - dE)e_i.$$

It is easy to see that the one-shot Nash equilibrium is such a setting entails emissions

$$e_i^N = \frac{2b_i - b_j}{3d}.$$

Combined emissions are $E^N = (b_1 + b_2)/3d$, so that Nash equilibrium payoffs for country i are

$$\pi_i^N = d(e_i^N)^2.$$

Both countries employ a discount factor δ to evaluate payoffs one period into the future. Mildly abusing notation, I denote the emissions selected by country $k = 1,2$ in period t as e_{kt}.

To support a cooperative regime (e_1^c, e_2^c) let us suppose the countries each play the grim strategy: country i chooses e_i^c in period 1; in any subsequent period $t > 1$, i chooses $e_{it} = e_i^c$ if $e_{ks} = e_k^c, k = 1,2$ in all previous periods $s < t$, otherwise choose $e_{it} = e_i^N$. With this strategy, there are two subgames of note: those where no player has deviated in any previous period, and those where one player has deviated in some previous period.[2] By design, the strategy dictates a

[2] If both players defected in the previous period, the typical convention is to treat such a period as if no defection occurred (Fudenberg and Tirole, 1991).

best-reply to the rival's strategy in the latter type of subgame (*i.e.*, choosing the Nash emissions is by definition a best-reply to the other country's Nash emissions), so the combination that obtains when both countries play the grim strategy induces a subgame-perfect Nash equilibrium when the strategy pair generates a Nash equilibrium. That in turn requires the present discount flow of payoffs associated with honoring the strategy, V_i^c, not be smaller than the present discounted flow of payoffs associated with defecting, V_i^d. The present discount flow of payoffs associated with honoring the strategy is $V_i^c = \frac{\pi_i^c}{1-\delta}$. Since defection will trigger reversion to the Nash equilibrium, and play will stay there forever after, the present discounted value of defection is easily seen to be $V_i^d = \pi_i^d + \frac{\delta}{1-\delta}\pi_i^N$, where π_i^d are the payoffs earned by selecting the one-shot best-reply to the rival firm's (cooperative) emissions. There are many combinations of emissions (e_1^c, e_2^c) that satisfy $V_i^c \geq V_i^d$; the boundary, where country i is just willing to play the grim strategy, is defined by $V_i^c = V_i^d$. In framework adopted here, this frontier is implicitly defined by a quadratic relation between e_i^c and e_j^c. I refer to this relation as the "incentive constraint" in the subsequent discussion.

Figure 1 illustrates the general principle. Here, I plot the incentive constraints for each of the two players. Choices for the larger player are plotted on the $x-$axis, while choices for the smaller player are plotted on the $y-$axis. The incentive constraints intersect at two places: the one-shot Nash equilibrium (the point farthest to the northeast) and the most cooperative outcome (the point farthest to the southwest). Also indicated in this diagram are combinations with the same ratio of emissions as in the Nash equilibrium, represented by the dashed line labeled "pro-rata sharing". The key point is that the pro-rata sharing locus crosses the large player's incentive constraint at a point well above the most cooperative regime, and so in general one might expect the large player to press for sharing rules that are disproportionately to its advantage. If, as seems intuitive, the smaller player insists on a more "equitable" sharing arrangement, it will be difficult to craft an agreement that exerts much influence on the levels of activity. In particular, it seems unlikely that a voluntary agreement will do much to reduce combined emissions.

3 The Experimental Data

To evaluate the predictions of the model I above, I make use of experimental data. In this experiment, subjects were placed in one of two structures, each of which was based on model described above. In both designs, $b_1 = 4$ and $d = \frac{1}{24}$. In the *symmetric* design, $b_2 = b_1$, while in the *asymmetric* design the "small" player has $b_2 = 3.5$. The one-shot Nash equilibrium choices are then $e_1^N = e_2^N = 32$ in the symmetric design and $e_1^N = 36, e_2^N = 24$ in the asymmetric design. Payoffs were presented to subjects in the form of payoff tables which show the payoff accruing from various output combinations.

Subjects were recruited for a length of time 30 to 45 minutes greater than an experimental actually ran. After the instructions were read, a practice period was conducted. A monitor randomly chose the counterpart value while all subjects simultaneously selected their row value from a sample payoff table. Then, half of the subject pool was moved to another room. Each person was matched with an anonymous opponent in the other room. Subjects were told they would be paired with the same person for the duration of the experiment. In each choice period the subjects wrote their choice on a record sheet and a colored piece of paper. These colored slips were then exchanged by a monitor, and payoffs for the period were tabulated from the payoff table. They had many more record sheets and colored slips of paper than required for a session. Each subject was given a starting cash balance of $5.00 to cover potential losses, and was told that if their balance went to zero they would be discussed from the experiment with a $2.00 participation fee (although this never happened). Subjects were told the experiment would run at least 35 periods, with a random termination rule (corresponding to a 20% chance that the experiment would be stopped) applied at the end of each period starting with period 35. Thus, the design mimics a repeated game with discount factor 0.8.

I use data from six experimental sessions. In three sessions based on the symmetric design, a total of 38 subjects (19 pairs) made choices for between 35 and 46 periods. In three asymmetric sessions, a total of 50 subjects (25 pairs) made choices for between for 36 to 46 periods. These data then induce an unbalanced panel; to avoid the potential for overly large influence on the results by subject pairs who were particularly fast at making decisions, or lucky in terms of the sequence of random termination actions, I restrict attention to the first 35 periods in the econometric analysis discussed below.

A visual summary of the experimental data is provided in Figures 2–4. Figure 2 shows the average choices made by subjects in the symmetric design, which I refer to as "LL." The Figure includes two reference levels of note: 32, the Nash equilibrium level, and 24, the joint payoff maximizing choice, Average choices tend to lie between these two reference levels, though closer to the Nash output. Figure 3 compares average market choices in the two designs. To facilitate comparison, I plot the choices as fractions of the market Nash equilibrium levels. The average market choices for the symmetric design are shown by the solid line, while the average choices for the asymmetric design are shown by the dashed line. On balance, the symmetric choices are a smaller fraction of the Nash equilibrium output. Figure 4 explores this latter pattern at the individual player level. Here, I plot the average choices for L (large) players as the solid line, with the average choices for S (small) players as the dashed line. The panel on the left shows the levels of choices, while the panel on the right shows these choices as fractions of the respective Nash equilibrium choices. That S player choices are smaller than L player choices, as depicted in the left panel, is to be expected in light of the payoff function disparity. But the interesting feature here is that S player choices are a larger fraction of the Nash equilibrium choice, on average, than are the L player choices. This pattern is at odds with the theoretical design articulated above, and indeed earlier conceptual analyses; this disparity merits deeper investigation, a task I undertake with a formal econometric analysis in the next section.

4 Econometric Analysis

The econometric model I employ treats the database as a pooled cross-section/time series sample. In this vein, I analyze choices made in each period for each of the subjects in a particular design, and analyze systematic differences in behavior to asymmetries in subjects' payoff functions.

I assume that an individual's choice in period t is related to the rival's choice in period t - 1, via a dynamic reaction function; this framework is similar to the empirical model discussed in Huck et al. (1999, eq. (4)). Because human subjects are likely to be boundedly rational, I allow for noise in this relation. Moreover, as there is a potential for learning or signaling (Mason and Phillips, 2001), the noise affecting the dynamic reaction function is likely to be serially correlated. Correspondingly, I allow the disturbance to follow an autoregressive process; in that way, the dynamic strategies can be rewritten as including N lags:

$$e_{it} = \varphi_{i0} + \Sigma_{n=1}^{N} \mu_{nh} e_{i,t-n} + \Sigma_{n=1}^{N} \nu_{nh} e_{j,t-n} + \omega_{kt} + \eta_{it}, \tag{1}$$

where e_{it} is player i's period t choice, j is $i's$ rival, k indexes the players' subject pair, $h = L$

(respectively, S) if player i is large (respectively, small), and I allow for individual-specific fixed effects (via φ_{i0}) and pair-specific variance (*i.e.*, random effects, via ω_{kt}^2). The individual-specific residual, η_{it}, is assumed to be white noise. I assume there is no cross-equation covariance between subject pairs. In the results reported below, I allow for $N = 3$ lags; with that structure the residuals display no serial correlation.[3]

I estimate the parameters in eq. (1) using random effects, including pair-specific dummy variables, while imposing robust standard errors (equivalently, clustering at the subject pair level). This approach yields consistent, asymptotically efficient estimates (Fomby et al., 1988).

Results from this regression analysis are collected in Table 1. Here I display parameter estimates for the asymmetric structure in regression 1 (the second column) and for the symmetric structure regression 2 in (the third column). To economize on space, I denote the explanatory variables in regression 2 as x_{nL} and $y_{nL}, n = 1,2,3$ (though all subjects in that treatment were type L). In both regressions, subjects tend to respond positively to their own past choices and negatively to their rival's past choices. In addition, small players in the asymmetric design, *i.e.*, subjects playing the role of small countries, choose markedly smaller values than do large players.

Once consistent and efficient estimates of the parameters are obtained, one can develop estimates of the underlying steady state (long run) equilibrium emission levels by considering the deterministic analogues to eq. (1). If agents choose the steady state values, e_L^* and e_S^*, for several

[3] The approach I took here was to collect the residuals e_{it} and then regress residuals at time t on residuals from time $t-1$; *i.e.*, $e_{it} = \rho e_{it-1} + u_{it}$. Observing a statistically important parameter estimate $\hat{\rho}$ indicates the presence of serial correlation. In the variant with $N = 3$ the parameter estimate $\hat{\rho}$ was not statistically significant ($\hat{\rho} = .178$; t-statistic = 1.44). I also estimated a variant of equation 1 with $N = 2$ lags; the residuals from that regression did display serial correlation ($\hat{\rho} = .196$; t-statistic = 7.99). I conclude from this exercise that the appropriate version of equation 1 has $N = 3$.

consecutive periods, this gives a system of two equations in two unknowns:

$$e_L^* = \varphi_{0L} + \mu_{1L}e_L^* + \mu_{2L}e_L^* + \mu_{3L}e_L^* + \nu_{1L}e_S^* + \nu_{2L}e_S^* + \nu_{3L}e_S^*, \tag{2}$$

$$e_S^* = \varphi_{0S} + \mu_{1S}e_S^* + \mu_{2S}e_S^* + \mu_{3S}e_S^* + \nu_{1S}e_L^* + \nu_{2S}e_L^* + \nu_{3S}e_L^*, \tag{3}$$

where φ_{0L} (respectively, φ_{0S}) refers to the average value of φ_{0i} across all L (respectively, S) subjects. Define $\tilde{\mu}_h = \mu_{1h} + \mu_{2h} + \mu_{3h}$ and $\tilde{\nu}_h = \nu_{1h} + \nu_{2h} + \nu_{3h}$, for $h = L, S$. Solving the system of equations (2)–(3) yields:

$$e_L^* = \frac{(1-\tilde{\nu}_S)\varphi_{0L} + \tilde{\nu}_L\varphi_{0S}}{(1-\tilde{\mu}_L)(1-\tilde{\nu}_S) - \tilde{\nu}_L\tilde{\mu}_S}, \tag{4}$$

$$e_S^* = \frac{\tilde{\mu}_S\varphi_{0L} + (1-\tilde{\mu}_L)\varphi_{0S}}{(1-\tilde{\mu}_L)(1-\tilde{\nu}_S) - \tilde{\nu}_L\tilde{\mu}_S}. \tag{5}$$

Inserting the estimates for the relevant parameters (taken from Table 1) into eqs. (4)–(5) then yields maximum likelihood estimates of the underlying equilibrium values.[4] Here, the resultant values are $e_L^* = 33.21$ and $e_S^* = 24.51$ for the asymmetric structure. These outputs are close to the one-shot Nash combination, indicating that subjects in the asymmetric structure were unable to effect much of a reduction in output. Moreover, most of the burden is carried by the larger player, as e_L^* is almost three units below player L's Nash choice, while e_S^* is slightly larger than S's Nash choice.

A similar approach may be used to estimate the equilibrium emissions choice in the sym-

[4] See Fomby et al. (1988) for details. Dynamic stability requires that all of the μ and ν parameters, as well as $1-\tilde{\mu}_h$ and $1-\tilde{\nu}_h, h = L, S$, are also less than one in magnitude – which they are here. This is a substantive concern, for dynamic stability allows one to interpret the carrot choices derived in eqs.eqs. (4)–(5) as equilibrium choices. Covariance information from the maximum likelihood estimates of the a's and b's can similarly be used to construct consistent estimates of the covariance structure for the steady state values (Fomby et al., 1988, Corollary 4.2.2).

metric structure. Here, however, the regression equation is

$$e_{it} = \varphi_{i0} + \Sigma_{n=1}^{N} \mu_n e_{i,t-n} + \Sigma_{n=1}^{N} \nu_n e_{j,t-n} + \omega_{kt} + \eta_{it}, \tag{6}$$

as all agents are type L. Accordingly, the steady state choice for subjects in the symmetric design is

$$e_L^* = \frac{\varphi_0}{1 - \tilde{\mu} - \tilde{\nu}}, \tag{7}$$

where as above $\tilde{\mu} = \mu_1 + \mu_2 + \mu_3$ and $\tilde{\nu} = \nu_1 + \nu_2 + \nu_3$, and where φ_0 refers to the average value of φ_{0i} across all subjects. Using the parameter values from Table 1, one obtains $e_L^{**} = 29.22$. Comparing this estimate against the estimate for e_L^*, I conclude that subjects in the symmetric structure were more successful at reducing emissions.

A clear implication of these results is that, at least in these experimental markets, it is difficult for the large player to induce the small player to act cooperatively.

5 Discussion

This paper highlights the importance of asymmetry in compromising cooperative cooperative behavior in climate negotiations. Both the analytics and the experimental evidence point to small player's behavior as key: While a cooperative regime would require tilting extraction in the direction of the larger player, small players resist. This intransigence ultimately undercuts the ability to form an arrangement that limits emissions.

One explanation for the results articulated above is that subjects' utility is based on both the payoffs they receive and the payoffs their rival receives. For example, if subjects bear disutility when there is disparity in payoffs, then subject i's utility can be expressed as

$$U_i(\pi_i, \pi_j) = \pi_i + \gamma|\pi_i - \pi_j|.$$

Denoting, as above, the large player as subject 1 and the small player as subject 2, and assuming $\pi_1 > \pi_2$, the two subjects' utilities can be written as

$$U_1 = (1-\gamma)\pi_1 + \gamma\pi_2;$$

$$U_2 = (1+\gamma)\pi_2 - \gamma\pi_1.$$

Grafting such a scenario onto the linear-quadratic framework introduced above, the one-shot Nash equilibrium would be

$$\tilde{q}_1^* = \frac{a+c-\gamma(a-c)}{3-4\gamma^2)b}; \tag{8}$$

$$\tilde{q}_2^* = \frac{a-2c+a\gamma}{3-4\gamma^2)b}. \tag{9}$$

It is easy to see that this pushes the Nash equilibrium towards a smaller (respectively, higher) output for the large (respectively, small) player. It also induces a similar effect on the quasi-cooperative player.[5] This adapted model also seems to coincide with casual empirical evidence taken from the historical record of international climate negotiations: developing countries long argued that developed countries should undertake larger emission reductions, often on the grounds

[5] An alternative explanation for the experimental outcome I describe in this paper is that small players are less patient, *i.e.* they use a smaller discount factor. As in the previous adaptation, this change will induce the high-cost (respectively, low-cost) player to select a larger (respectively, smaller) output in the quasi-cooperative arrangement than is predicted in the model with a common discount factor.

of equity.

References

Barrett, S. (1994). Self-enforcing international environmental agreements, *Oxford Economic Papers* **46**: 878–894.

Barrett, S. (2008). Climate treaties and the imperative of enforcement, *Oxford Review of Economic Policy* **24**: 239–258.

Bolton, G. E. and Ockenfels, A. (2000). ERC: A Theory of Equity, Reciprocity, and Competition, *The American Economic Review* **90**: 166–193.

Carraro, C. and Siniscalco, D. (1993). Strategies for the international protection of the environment, *Journal of Public Economics* **52**: 309–328.

Finus, M. (2001). *Game Theory and International Environmental Cooperation*, Edward Elgar, London.

Fomby, T. B., Hill, R. C., and Johnson, S. R. (1988). *Advanced Econometric Methods*, 2nd edn, Springer-Verlag, New York.

Fudenberg, D. and Tirole, J. (1991). *Game Theory*, MIT Press, Cambridge, MA.

Huck, S., Normann, H.-T. and Oechssler, J. (1999). Learning in Cournot Oligopoly – An Experiment, *The Economic Journal* **109**: C80 – C95.

Mason, C. F. and Phillips, O. R. (2001). Dynamic learning in a two-person experimental game, *Journal of Economic Dynamics and Control* **25**: 1305–1344.

Schmalensee, R. (1987). Competitive Advantage and Collusive Optima, *International Journal of Industrial Organization* **5**: 351–367.

Figure 1: Asymmetric Incentives to Cooperate.

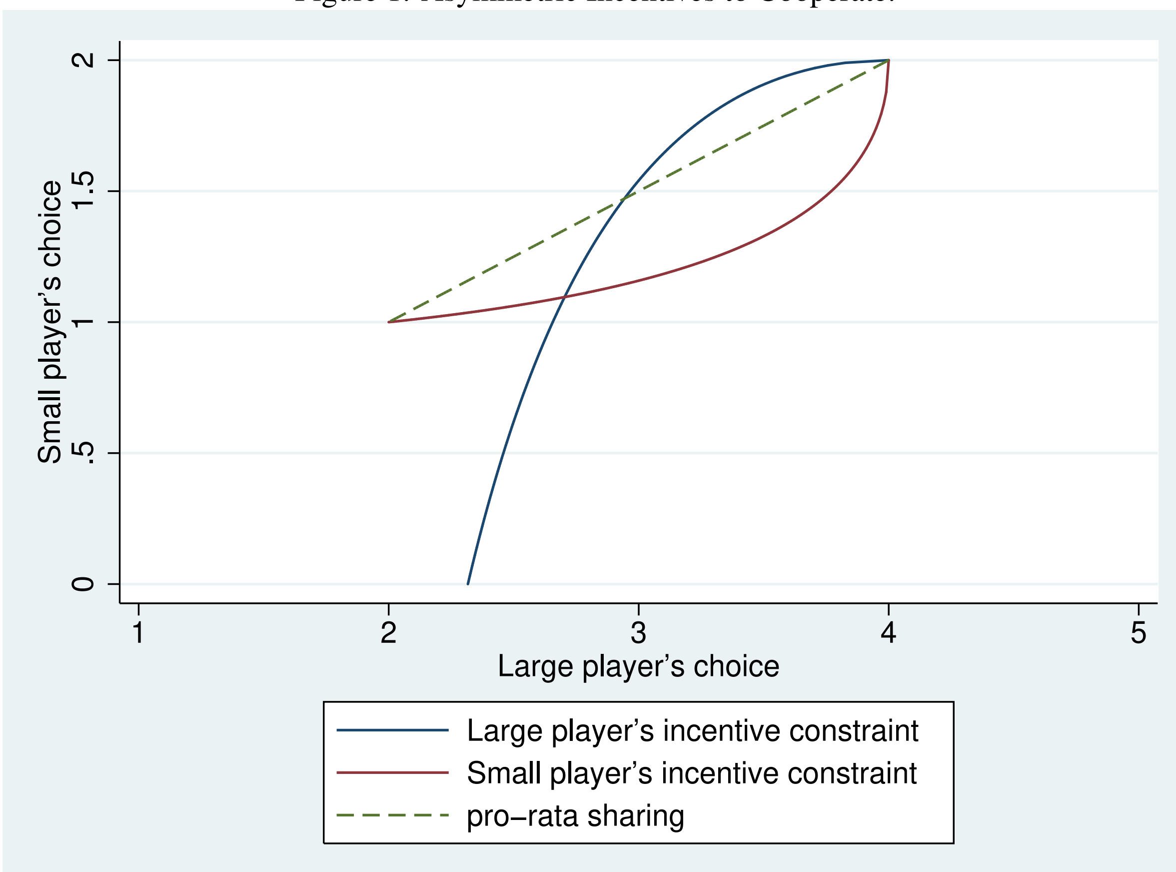

Figure 2: Experimental results: symmetric firms.

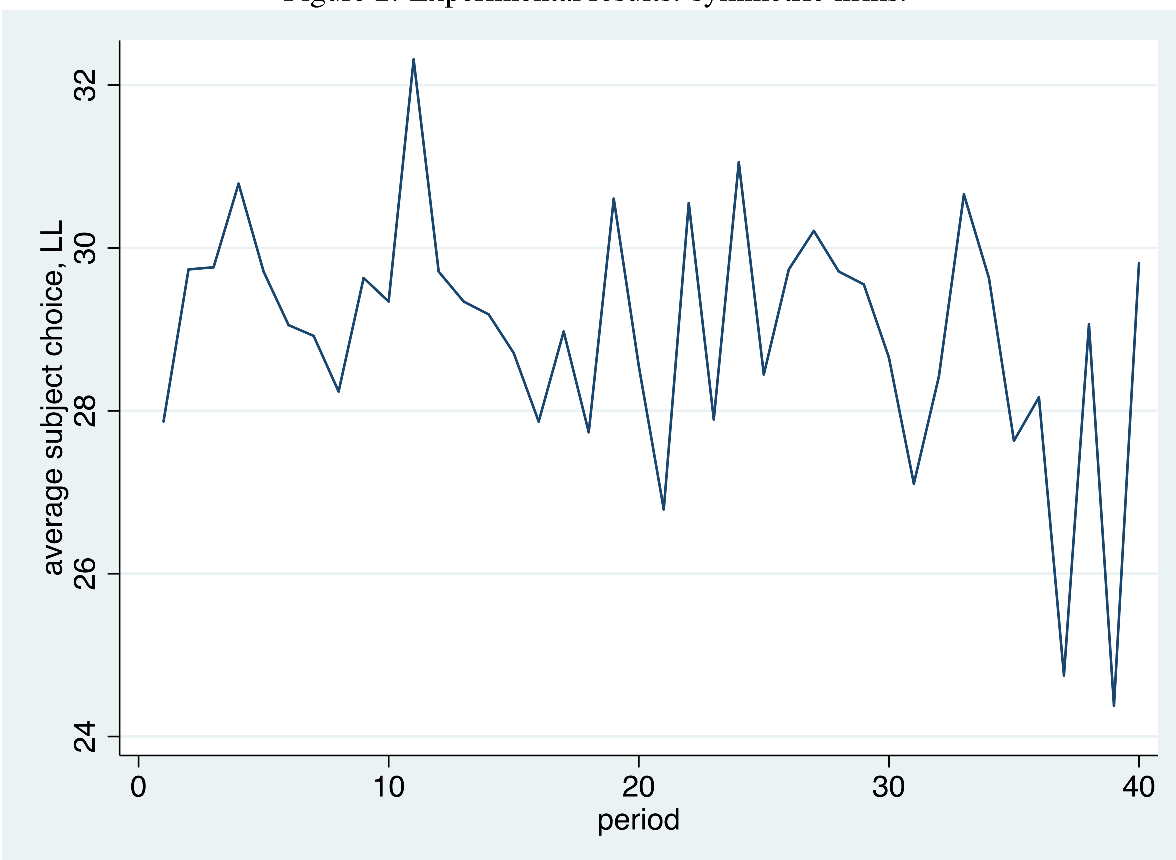

Figure 3: Experimental results: symmetric vs. asymmetric designs.

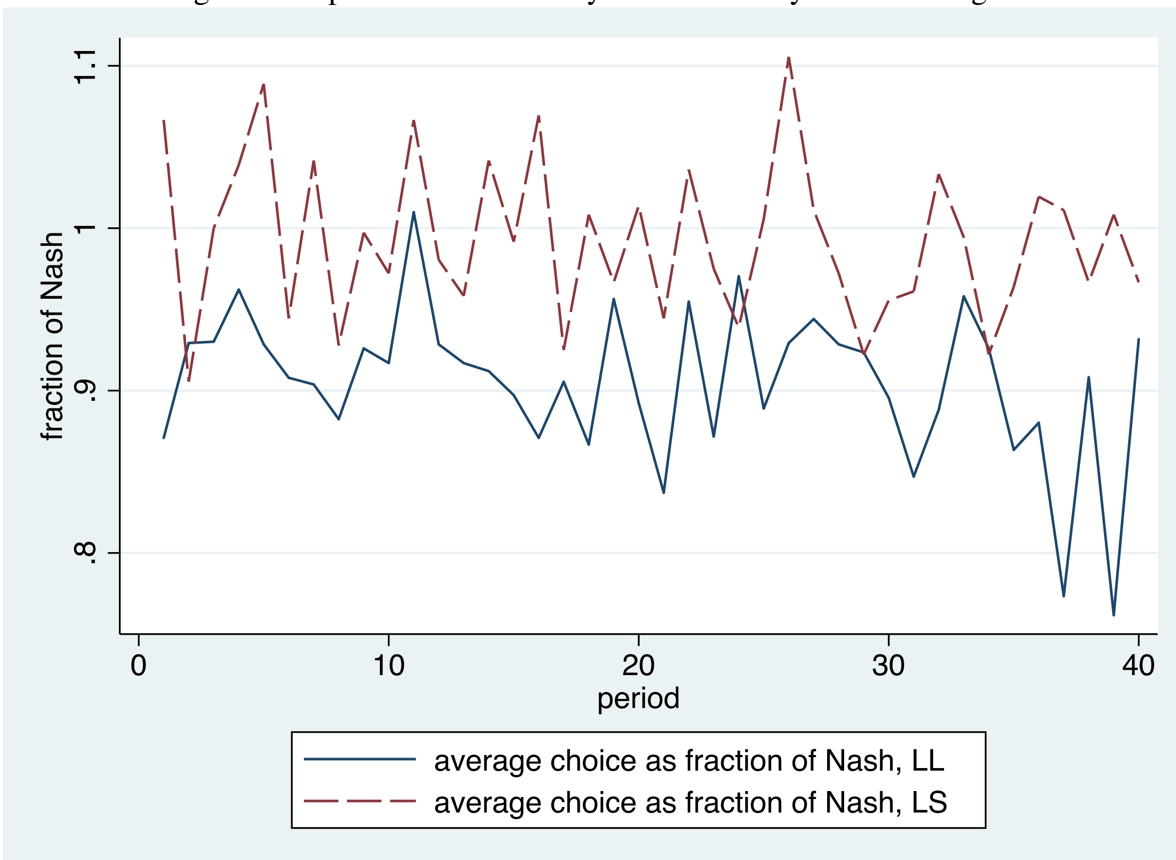

Figure 4: Experimental results: asymmetric countries.

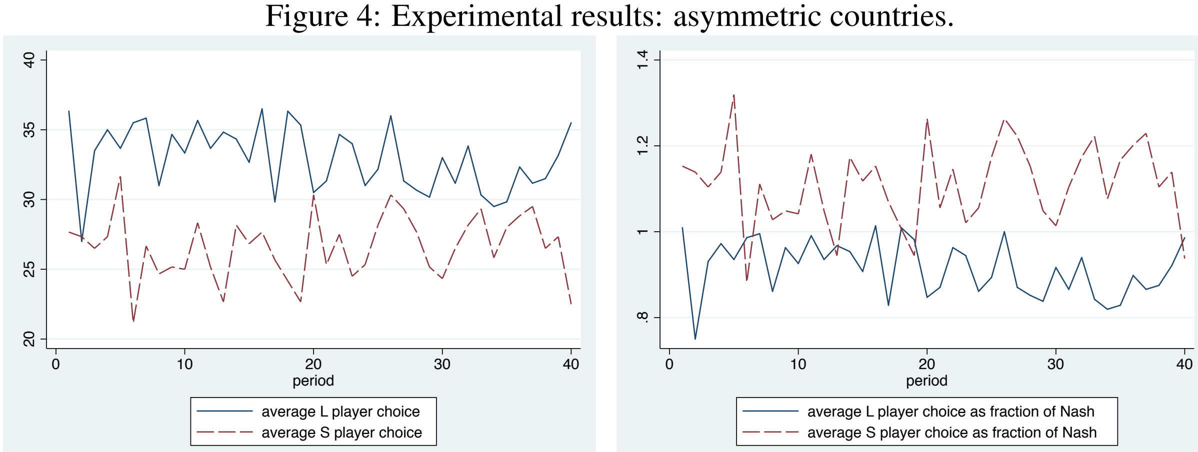

Table 1: Random Effects Regression Results: Asymmetric vs. Symmetric Experimental Designs.

	asymmetric	symmetric
xlL	-0.264***	-0.327***
	(0.069)	(0.074)
xl2L	0.151	0.026
	(0.107)	(0.072)
xl3L	-0.040	-0.131***
	(0.050)	(0.030)
ylL	0.253**	0.210***
	(0.099)	(0.060)
yl2L	-0.033	-0.022
	(0.070)	(0.040)
yl3L	0.075	0.103**
	(0.066)	(0.046)
xlS	-0.050	
	(0.055)	
xl2S	0.100	
	(0.141)	
xl3S	-0.059	
	(0.064)	
ylS	0.182***	
	(0.065)	
yl2S	-0.141***	
	(0.048)	
yl3S	0.029	
	(0.048)	
constant	21.064***	42.757***
	(4.233)	(4.300)
$Q_{L}*$	33.21	29.22
	(1.986)	(0.945)
$Q_{S}*$	24.51	—
	(4.312)	
N	1600	1216
R^2	0.5375	0.4350

All regressions include individual-specific dummy variables

Robust standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$